Technical white paper

HP RA for Red Hat Storage Server on HP ProLiant SL4540 Gen8 Server

Deploying open scalable software-defined storage

Table of contents

Executive summary

Organizations of all kinds are struggling with the need to manage growing amounts of data in a cost-effective fashion, even as they focus on extracting more value from that data. Unstructured data in particular continues to grow rapidly in the typical enterprise, outpacing the abilities of conventional storage solutions. Enterprises have seen enormous gains in scalability, flexibility, and affordability as they migrated from proprietary, monolithic server architectures, but storage has remained a bastion of inflexible technology. The advent of virtualization and growing adoption of cloud models are only adding to these challenges.

Driven by the need for improved scalability, performance, and faster provisioning at lower infrastructure and management costs, HP and Red Hat® are collaborating to bring effective scale-out storage solutions to market. Software-defined storage solutions provide a compelling alternative for serving growing amounts of unstructured data in an approach that scales easily in multiple dimensions. Scale-out storage increases performance, capacity, or throughput by adding resources (processors, memory, network interfaces, or spindles) as a loosely coupled system of nodes that work side-by-side, in parallel.

Red Hat Storage Server running on HP ProLiant SL4540 Gen8 servers offers massive scalability, high performance, and volume economics for storage of all kinds of unstructured data. In particular, two-node HP ProLiant SL4540 Gen8 servers equipped with up to 25 disks per node can offer scalable and cost-effective solutions for a variety of use cases, including:

- Creating a large file and object store
- Deploying an enterprise drop box or cloud storage for service providers
- Establishing a near-line archive

Target audience: Technical decision makers, datacenter managers, and storage managers wishing to learn more about deploying scale-out storage using Red Hat Storage Server and HP ProLiant SL4540 Gen8 servers.

Document purpose: The purpose of this document is to describe a recommended architecture/solution, highlighting recognizable benefits to technical audiences.

This white paper describes testing performed by Red Hat and HP in January 2014.

Introduction

Red Hat Storage Server offers software-only scale-out storage, which, when combined with the HP ProLiant SL4540 Gen8 Server, provides a unique storage solution that delivers innovative functionality. The all-inclusive Red Hat Storage Server license includes such features as Monitoring, Tiering, Quotas, Load balancing and failover, Geo Replication and Snapshots (available June 2014), all of which may incur additional license fees with other solutions. Along with the flexibility of the HP ProLiant SL4540 Gen8 Server, HP and Red Hat can deliver a cost-effective scale-out storage solution to meet different workload requirements, whether offering the high capacity requirements of near-line archiving, or high performance for online analytics.

Storage growth, innovation, and cost

In recent years, unstructured data has been growing at unprecedented rates. Enterprises of all kinds are facing an explosion of data from diverse applications and technology trends including virtualization, collaboration, business intelligence, data warehousing, and media. The proliferation of sensors and mobile devices along with the move to cloud computing models will only accelerate these growth trends. New regulation and data retention requirements further exacerbate the problem. In fact, according to IDC, from 2005 to 2020, the digital universe will grow by a factor of 300, from 130 exabytes to 40,000 exabytes[1] with 30.1 billion connected devices by 2020[2]. At the same time, managing data and exploiting it for new opportunities has become a top priority for IT organizations as the strategic and tactical importance of data and its analysis have become apparent.

Driven by this rampant growth, storage costs are having an outsized impact on enterprise IT budgets. While the volume and use of storage is expected to grow dramatically each year, the price for storage continues to remain high. Per-terabyte storage prices are not falling as fast as data volumes are growing, presenting a problematic diverging trend for organizations that must supply and manage large-scale storage infrastructure. For example, IDC estimates that the investment in spending on IT hardware, software, and services, telecommunications, and staff will grow by 40% between 2012 and 2020.[1]

[1] Source: IDC, The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East
[2] Source: IDC Press Release, Oct 3, 2013: “The Internet of Things Is Poised to Change Everything”

As a result of these trends, organizations are frequently forced into buying expensive solutions to meet their storage requirements. Every time organizations want to add more data, they have to purchase additional storage capacity. Every time they add more capacity, they have to sacrifice performance, as many storage solutions do not scale linearly. Making matters worse, many storage solutions often have rigid architectures with limited deployment options, and are not flexible enough to both scale up, by adding more capacity, and scale out, by adding more nodes. Furthermore, JBOD (just a bunch of disk) solutions frequently introduce complexity through considerable hardware variety, which lacks a cohesive scalability strategy and management model. Strategies are needed that enable both cost-effective growth and scalability of storage for unstructured data.

Scale-out storage with HP and Red Hat

Once thought of as a niche approach, software defined storage has emerged as a compelling solution. Just as scale-out server architectures have revolutionized cost-effective computing, scale-out storage offers many of the same benefits for storage. Virtualization and cloud computing benefit in particular, with software defined storage offering a more scalable storage infrastructure to support rapid virtual machine growth. The same aspects of scale-out storage that make it ideal for virtualization and cloud make it ideal for serving unstructured data of all kinds.

Red Hat Storage Server running on HP ProLiant SL4540 Gen8 servers represents a purpose-built solution that takes scale-out storage to a new level of density and volume economics—allowing organizations to effectively store and manage petabyte-scale data growth. Together, Red Hat Storage Server and the flexible and scalable HP ProLiant SL4540 Gen8 server offer cost-effective multi-dimensional storage scalability. In fact, Red Hat and HP estimate that a 5PB installation based on Red Hat Storage Server and HP ProLiant SL4540 Gen8 servers could offer a 30% savings in infrastructure, space, power, and cooling costs over a comparable JBOD configuration over a three-year period, including:

- 50% less space
- 61% less power
- 75% less administration time
- 63% fewer cables

The HP ProLiant SL4540 Gen8 Server is an excellent choice for running Red Hat Storage Server and is available in one of three configurations (each of these configurations is available in a 4.3U form factor chassis):

- 1 node configuration – one compute node with up to two small form factor (SFF) 2.5" hard disk drives (HDD) and up to sixty (60) large form factor (LFF) 3.5" hard disk drives (HDD)
- 2 node configuration – two compute nodes, each with up to two small form factor (SFF) 2.5" hard disk drives (HDD) and up to twenty-five (25) large form factor (LFF) 3.5" hard disk drives (HDD)
- 3 node configuration – three compute nodes, each with up to two small form factor (SFF) 2.5" hard disk drives (HDD) and up to fifteen (15) large form factor (LFF) 3.5" hard disk drives (HDD)

This document will focus on the 2 node configuration of the HP ProLiant SL4540 that can hold a total of fifty (50) large form factor (LFF) 3.5" hard disk drives (HDD) in the storage section and four small form factor (SFF) 2.5" HDDs (two for each server node). Each node within the HP ProLiant SL4540 is a dual socket Gen8 server, with a choice of Intel® Xeon® processors ranging from quad-core to ten-core processors, up to 192GB of memory and a single PCIe slot for expansion. Furthermore, you have a choice of mixing SATA, SAS, and SSD drives within the chassis depending on the disk performance needed for your application, and there is an option for Flash IO Accelerator cards to enhance IO capabilities.

Overview

The combination of Red Hat Storage Server and HP ProLiant SL4540 Gen8 servers provides a next-generation scale-out storage solution. Red Hat Storage Server is a software-defined scale-out file and object storage software solution for private cloud or datacenter, public cloud, and hybrid cloud environments. Red Hat Storage Server is software-only, open source, and designed to meet unstructured, semi-structured, and big data storage requirements.

Red Hat Storage Server enables organizations to combine large numbers of HP ProLiant SL4540 Gen8 servers into a high-performance, virtualized, and centrally-managed storage pool. Both capacity and performance can scale linearly and independently on demand, from a few terabytes to petabytes and beyond. By combining economical HP ProLiant SL4540 Gen8 servers with a scale-out approach, organizations can achieve radically better price and performance in an easily deployed and managed solution that can be configured for increasingly demanding workloads.

Benefits of the joint Red Hat and HP solution include:

- Flexible deployment. Flexible configurations of compute and storage can be tuned for diverse workloads.
- Drastically lower cost. Massive scalability and density provided by HP ProLiant SL4540 Gen8 servers consolidate power and cooling and save on space.
- Unstructured data scalability. Scale-out storage software offers a scalable foundation for unstructured data.
- Open and extensible. Extensible storage modules enable new business models without concern for inflexible infrastructure.

The HP and Red Hat solution provides multi-dimensional scale-out storage infrastructure in terms of both performance and capacity. By aggregating the disk, CPU, and I/O resources of large numbers of HP ProLiant SL4540 Gen8 servers, organizations can create large high-performance storage pools. More capacity can be added simply by adding additional disks. More performance can be added by adding additional server nodes (Figure 1). The unique architecture of Red Hat Storage Server delivers the benefits of scale-out storage while avoiding the corresponding overhead and risk associated with keeping large numbers of storage nodes in sync. In practice, Red Hat Storage Server accomplishes multi-dimensional scalability through:

- The elimination of the metadata server
- Effective distribution of data to achieve scalability and reliability
- The use of parallelism to maximize performance via a fully distributed architecture

Figure 1. Red Hat Storage Server running on the HP ProLiant SL4540 Gen8 server offers both performance and capacity scalability.

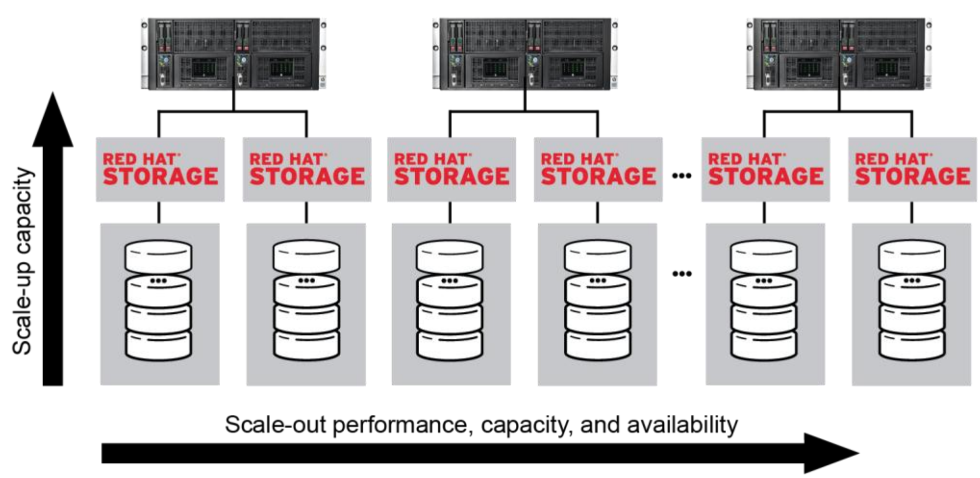

The joint HP and Red Hat solution delivers a number of benefits, including:

- Reduced costs. Deploying scale-out storage with Red Hat Storage Server on HP ProLiant SL4540 Gen8 servers allows organizations to dramatically reduce capital costs while maintaining high levels of performance and availability. Furthermore, all of the features offered by Red Hat Storage Server are included in one low license fee.
- Outstanding flexibility. The HP ProLiant SL4540 Gen8 Server can be configured with one, two, or three compute nodes, each with two Intel processor sockets supporting several processor choices, depending on your core to spindle performance requirements. You can also combine SATA, SAS, and SSD drives within each storage tray, giving you maximum flexibility in your disk performance. Each node also has the option of supporting Flash IO Accelerators for additional performance, and the cluster can be networked using 1Gb or 10Gb Ethernet, or via InfiniBand, depending on the network traffic requirements. All of this offers the flexibility to meet the performance needs of the workload while maintaining a cost-effective solution.
- Enhanced agility. Red Hat Storage Server can be deployed and scaled in minutes. Because Red Hat Storage Server automates the management of files and storage nodes, operational complexity is simplified and costs are reduced.
- Increased reliability and reduced downtime. Software ensures the availability of the storage system and its data. Red Hat Storage Server can replicate data to multiple HP ProLiant SL4540 Gen8 servers, helping to ensure that the system is protected from faults and that the failure of any individual server does not compromise data access, or the overall availability of the storage system.

- Lower complexity. Red Hat Storage Server running on HP ProLiant SL4540 Gen8 servers provides vastly simplified homogeneous infrastructure over disparate JBOD systems, while a simplified global namespace helps eliminate architectural complexity.
- Better user and customer satisfaction. By consistently delivering high levels of performance and helping to eliminate data silos, Red Hat Storage Server combined with HP ProLiant SL4540 Gen8 servers helps drive higher user satisfaction levels.

The sections that follow describe how the combination of Red Hat Storage Server with the HP ProLiant SL4540 Gen8 server can be deployed to serve a diverse range of use case scenarios.

Large file and object store

Many organizations use applications that generate a vast amount of unstructured data. Taking the form of either files or objects or both, these data need to be stored for analysis and synthesis at a later time. Industries as diverse as oil and gas-exploration (seismic data), telecommunications (data and application logs and billing records), healthcare (imaging and patient records), finance (risk analysis, compliance, and security data), scientific (gene sequencing and simulations) and others (rendering, geographical information systems and computer aided engineering) all generate vast amounts of data. Unfortunately, these unstructured data can accumulate at an exponential rate, causing organizations to continue to react by adding additional file servers or expensive monolithic NAS devices. Over time, these sprawling configurations can become difficult to manage and control. In many cases, this file server sprawl can result alternately in underutilization of costly storage devices or lower performance because of over-utilization of those same file servers and NAS devices.

Large file and object store workloads can be characterized as write once (or write-seldom) with many reads. Average file sizes can vary from small to large, or from a few kilobytes to several megabytes. The total size of the data store can start small (from a few terabytes) but will often grow to hundreds of terabytes and continue to grow steadily. The access pattern generally consists of several concurrent users reading and only a few writing at any given point in time.

Figure 2 illustrates multiple HP ProLiant SL4540 Gen8 servers combined via Red Hat Storage Server into a single volume that serves as a large file and object store. This approach provides a number of advantages when serving as a large file and object store, including:

- Elastic scalability
- Elimination of storage silos
- Centralized and simplified management with global namespace technology
- Volume server and storage economics
- Multi-protocol client support for flexible file sharing

Figure 2. Red Hat Storage Server and the HP ProLiant SL4540 Gen8 server can be deployed to serve as a large file and object store.

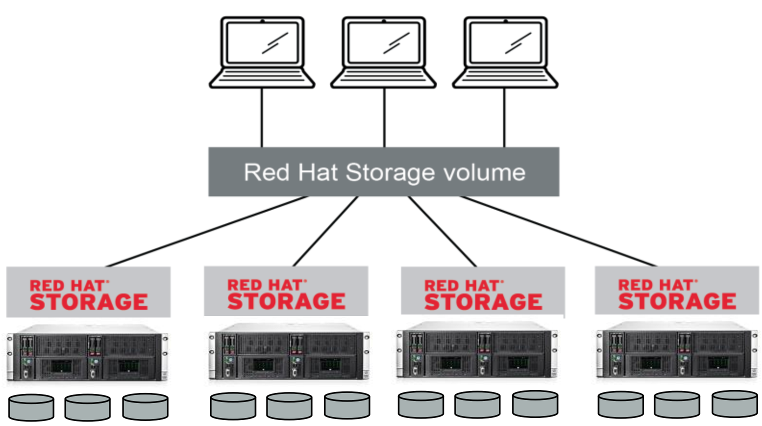

Enterprise drop box or cloud storage for service providers

Telecommunication companies, cable operators, and other content providers are increasingly faced with competitive needs for providing differentiated services for their end customers. As a consequence, many service providers offer their customers access to free and/or paid storage for storing photos, videos, and other media files that they access via the Web using a smart phone, tablet, or a personal computer. A similar use case is seen in the enterprise context, where employees need to reliably and securely store and retrieve files from a variety of devices. Many find that building a reliable back-end storage infrastructure to meet the agility, cost-point, and versatility of demanding cloud storage or enterprise drop box deployments is extremely challenging.

The enterprise drop box or cloud storage workload can be characterized as initially comprising only writes, as end-users upload media files. Over time reads and writes become roughly equivalent as users access, view, and play those media files. File sizes range from a few megabytes (smart phone photos) to dozens of megabytes (HD videos taken using high-end digital SLR cameras). The total size of the data store and its growth is directly proportional to the number of users of the service. Deployments typically start at hundreds of terabytes and there is a need to pre-provision for potentially millions of subscribers or enterprise users. As new users sign up for content services, typical data stores can easily grow to several petabytes over time, and often span multiple datacenters. Data access patterns consist of hundreds of concurrent users reading and writing at the same time.

Competitive storage solutions must be able to deliver very high sustained throughput. As shown in Figure 3, multiple HP ProLiant SL4540 Gen8 servers are clustered via Red Hat Storage Server to provide a unified namespace, allowing access from diverse devices in either an enterprise drop box or cloud storage setting.

Together Red Hat Storage Server running on the HP ProLiant SL4540 Gen8 server provides a scale-out architecture with distinct advantages that include:

- Massive and linear scalability, with capacity and performance that can scale independently with no bottlenecks
- Object storage support via the OpenStack SWIFT API
- Support for multiple drop box software implementations that can be co-resident on the storage servers
- Simultaneous storage and retrieval of files and objects interchangeably
- Consistent, sustained, and predictable throughput performance

Figure 3. Red Hat Storage Server and HP ProLiant SL4540 Gen8 servers provide a unified namespace and can be deployed to serve as an enterprise drop box, or as cloud storage for service providers.

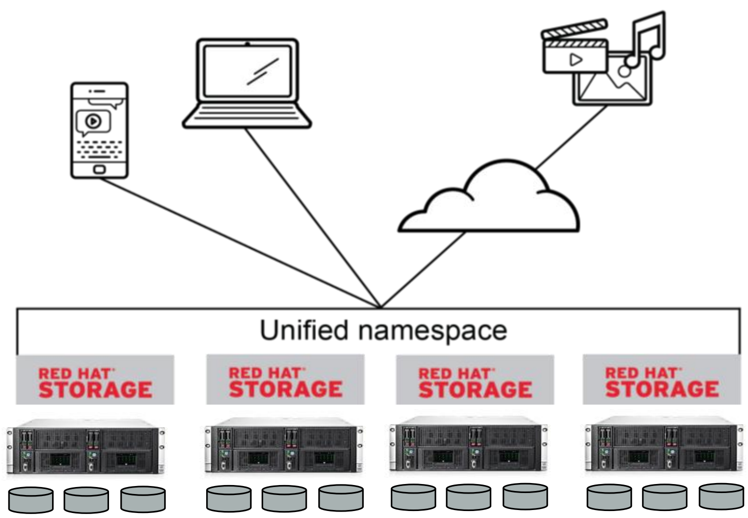

Near-line archival

A wide variety of organizations need to backup vast amounts of unstructured data, quickly accessing archived data when needed. Financial services firms, pharmaceutical companies, logistics operations, telecommunications, oil and gas, healthcare, public safety, educational, and government organizations all share this common need. Complicating matters, new regulations now dictate that organizations must comply with multi-year data retention policies. Customers are also demanding stringent Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO).

Near-line archival strategies often involve replacing tape archival with faster disk based access. Consistent access to data regardless of where it resides with consistent performance is another key requirement. Many storage solutions are simply unable to keep pace with the explosive growth of data and information.

The near-line archival workload can be characterized as write once, read sparingly or never. The average file sizes depend on the kind of data that is being backed up by the organization. The total data store usually starts with dozens of terabytes and can easily grow to several petabytes. However, if tape replacement is involved, the data store may start at a petabyte scale. The data access pattern consists of concurrent writes as the archive is populated, and typically very few if any reads.

As shown in Figure 4, multiple Red Hat Storage Server volumes can be deployed with one serving as near-line archival to the other. For example, a Red Hat Storage Server-based object store volume or a drop box volume could use another Red Hat Storage Server-based volume for near-line archival. The HP and Red Hat solution provides distinct advantages, including:

- Elastic scalability
- Rapid and random file access to ensure speedy access and recovery (faster retrieval and restoration of file data)
- Optional high availability via replication
- Optimized costs for archiving unprecedented amounts of data while minimizing complexities
- Compatibility with industry standards, including support for CIFS, NFS, and HTTP protocols as well as POSIX compatibility
- Reduced backup windows

Figure 4. Red Hat Storage Server running on HP ProLiant SL4540 Gen8 servers can function as an ideal near-line archival solution.

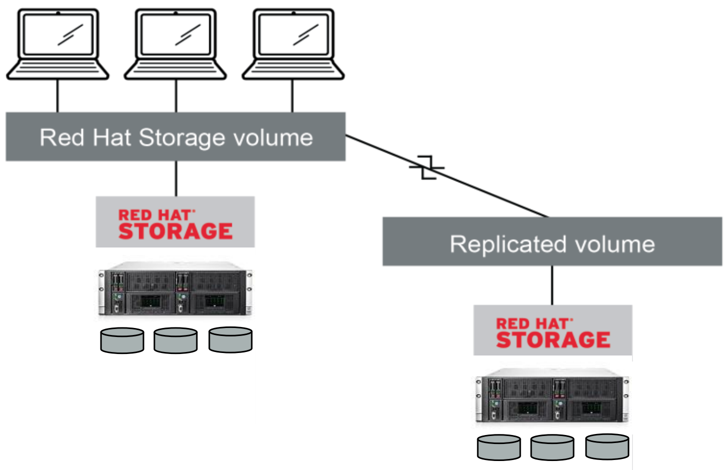

Solution components

Red Hat Storage Server and the HP ProLiant SL4540 Gen8 server represent an ideal technology combination for deploying scale-out storage.

Red Hat Storage Server

Red Hat Storage Server is software-only, open source, and designed to meet unstructured, semi-structured, and big data storage requirements. At the heart of Red Hat Storage Server is GlusterFS, an open source, massively scalable distributed file system that allows organizations to combine large numbers of storage and compute resources into a high-performance, virtualized, and centrally-managed storage pool. Red Hat Storage Server was designed to achieve several major goals as described in the sections that follow.

- Elasticity. With Red Hat Storage Server, storage volumes are abstracted from the hardware with each managed independently. Volumes can grow or shrink by adding or removing systems from the storage pool. Even as volumes change, data remains available with no application interruption.
- Petabyte scalability. Today's organizations demand scalability from terabytes to multiple petabytes. Red Hat Storage Server lets organizations start small and grow to support multi-petabyte repositories as needed. Organizations that need very large amounts of storage can deploy massive scale-out storage from the outset.
- High performance. Red Hat Storage Server provides fast file access by eliminating the typical centralized metadata server. Files are spread evenly throughout the system, eliminating hot spots, I/O bottlenecks, and high latency. With the HP ProLiant SL4540 Gen8 server, organizations can use commodity disk drives and 10 Gigabit Ethernet to maximize performance.
- Reliability and high availability. Red Hat Storage Server provides automatic replication that helps ensure high levels of data protection and resiliency. In addition to protection from hardware failures, self-healing capabilities restore data to the correct state following recovery.
- Industry standard compatibility. For any storage system to be useful, it must support a broad range of file formats. Red Hat Storage Server provides native POSIX compatibility as well as support for CIFS, NFS, and HTTP protocols. The software is readily supported by off-the-shelf storage management software.

- Unified global namespace. Red Hat Storage Server aggregates disk and memory resources into a single common pool. This flexible approach simplifies management of the storage environment and eliminates data silos. Global namespaces may be grown and shrunk dynamically, with no interruption to client access.
- Rapid and random access. Unlike archival systems based on tape, Red Hat Storage Server offers fast random access, and data can be both accessed and recovered rapidly.

From a technical perspective, Red Hat Storage Server provides distinct advantages over other technologies.

- Software-defined storage. Red Hat believes that storage is a software problem that cannot be solved by locking organizations into a particular storage hardware vendor or a particular hardware configuration. Instead, Red Hat Storage Server is designed to work with a wide variety of industry-standard storage, networking, and compute server solutions.
- Open source. Red Hat believes that the best way to deliver functionality is by embracing the open source model. As a result, Red Hat Storage Server users benefit from a worldwide community of thousands of developers who are constantly testing the product in a wide range of environments and workloads, providing continuous feedback, and providing unbiased feedback to other users.
- Complete storage operating system stack. Red Hat Storage Server delivers more than just a distributed file system. The complete storage solution adds distributed memory management, I/O scheduling, software RAID, self-healing, local N-way synchronous replication, and asynchronous long-distance replication via Red Hat Geo-Replication.
- User space. Unlike traditional file systems, Red Hat Storage Server operates in user space, rather than kernel space. This innovation makes installing and upgrading Red Hat Storage Server significantly easier, and greatly simplifies development efforts since specialized kernel experience is not required.
- Modular, stackable architecture. Red Hat Storage Server is designed using a modular and stackable architecture approach. Configuring Red Hat Storage Server for highly-specialized environments is a simple matter of including or excluding particular modules.
- Data stored in native formats. With Red Hat Storage Server, data is stored on disk using native formats (i.e., XFS) with various self-healing processes established for data. As a result, the system is extremely resilient and files are naturally readable without Red Hat Storage Server.
- No metadata with the elastic hash algorithm. Unlike other storage systems with a distributed file system, Red Hat Storage Server does not create, store, or use a separate index of metadata in any way. Instead, Red Hat Storage Server places and locates files algorithmically. The performance, availability, and stability advantages of not using metadata are significant, and in some cases, dramatic.

Two-node HP ProLiant SL4540 Gen8 server

The HP ProLiant SL4500 Gen8 server is the industry's first ever purpose-built server for big data. With support for up to 540 drives and over 2PB per rack, these servers are ideal for building dense scale-out storage infrastructure. At the same time, the server features a converged and balanced architecture that is ideal for a range of workloads, with an efficient chassis that scales easily as data grows. The server is designed with workload optimization in mind, with different ratios of compute and storage capacity per server, and with scalable performance and throughput demanded by those workloads.

- The one-node chassis supports up to 60 disk drives.
- The two-node chassis supports up to 25 disk drives per node (up to 50 drives per chassis).
- The three-node chassis supports up to 15 disk drives per node (up to 45 drives per chassis).

The HP and Red Hat scale-out storage solution and use cases described in this document are based on the two-node HP ProLiant SL4540 Gen8 server. However, depending on application needs, Red Hat Storage Server can be run on one-, two-, and three-node HP ProLiant SL4540 Gen8 servers as described below in the section on capacity and sizing. Importantly, each node within the two-node and three-node chassis is electrically isolated. Power and cooling are shared, reducing both CAPEX and OPEX required for the entire system. HP ProLiant SL4540 Gen8 servers can be configured to order.

For all of the HP ProLiant SL4540 Gen8 servers, each compute node is further divided into three zones to provide serviceability and balanced performance:

- Compute node module. The compute module includes two processor sockets and 12 DIMM slots. One or two Intel E5-2400 or E5-2400 v2 series processors can be configured. Two 2.5-inch small form factor (SFF) hot pluggable SATA drive slots provide redundant boot and operating system space, allowing all of the large form factor (LFF) drives to be used for data storage.
- I/O module. The I/O module includes an HP Smart Array P420i controller and an optional low-profile PCI Express slot. Dual Gigabit Ethernet connections are provided and an optional dual-port 10 Gigabit Ethernet NIC is supported (upgradable to a single 40Gb InfiniBand port).

- Capacity area. The capacity area for each server includes space for 15, 25, or 60 LFF hard disk drives, depending on the model chosen. SATA and SAS 7.2K rpm drives are supported as well as SAS 15K rpm drives. Up to 4TB drives can be configured. Solid state devices (SSDs) are also available in an LFF carrier.

HP ProLiant SL4540 servers were designed to be a part of enterprise-class solutions. As such, these servers offer key enterprise features. Cluster management is supported using the HP Insight Cluster Management Utility. Sophisticated power management is provided including Dynamic Power Capping and Asset Management with SL Advanced Power Manager. 94% Platinum Plus efficient power supplies offer power discovery. A host of additional enterprise-class features make the HP ProLiant SL4540 ideal for large-scale deployments in both enterprise settings and cloud service provider environments.

- Serviceability. In complex datacenters, scale-out data deployments need each node to be serviceable, and service times need to be quick. With HP ProLiant SL4540 Gen8 servers, compute nodes are serviceable from the front (hot aisle) side of the chassis. No cables need to be disconnected from the rear of the unit to service motherboard components. The I/O module is serviceable from the rear of the chassis, and the I/O module and PCI cards can be swapped without having to walk to the front of the unit. No cables cross a service zone, and the HDD backplanes are cable-less (and can be rapidly swapped as well).
- Performance. Matching processor cores to spindles is one aspect of performance that is important for being able to configure workload-specific server infrastructure. Another important aspect is the relationship of network performance to compute or storage performance. HP ProLiant SL4540 Gen8 servers provide a truly balanced architecture, with 40Gbps of raw bandwidth on the network side and 48Gbps of raw bandwidth on the SAS controller side. Network bandwidth requirements vary, so various I/O options are provided including Gigabit Ethernet, 10 Gigabit Ethernet, and InfiniBand compatibility.
- HP SmartCache technology. HP SmartCache technology accelerates access times to most frequently-accessed data (hot data). With support provided on the HP Smart Array P420i controller, HP SmartCache technology caches hot data on solid-state devices (SSDs). This approach combines low cost, high capacity spinning media with fast, low-latency SSD devices to dramatically accelerate performance over spinning media alone, at a fraction of the cost of an end-to-end solid-state solution.[3]

Figure 5. HP ProLiant SL4540 Gen8 Server (2 node configuration with top cover removed)

[3] Red Hat Storage Server has not yet been tested with SSDs or HP SmartCache technology as of this writing. Future testing is under investigation.

Switches

Red Hat Storage Server clusters can be connected using 1Gb Ethernet or 10Gb Ethernet ports. In our testing, we used an HP 5920 series switch (Figure 6) to cluster nodes via 10Gb Ethernet for fast performance.

Top of Rack (ToR) switches (5920AF-24XG Switch)

The HP 5920AF-24XG 10GbE high-density, ultra-deep packet buffering, top-of-rack (ToR) switch provides IRF Bonding and sFlow which simplifies the management, monitoring and resiliency of a Red Hat Storage Server network. This model has 24 x 10-Gigabit / Gigabit SFP+ ports for ultra-high capacity connections. The high performance 10GbE networking provides cut-through and nonblocking architecture that delivers industry-leading low latency (~1.7 microsecond) for very demanding enterprise applications; the switch delivers a 480Gbps switching capacity and 357.12 Mpps packet forwarding rate in addition to incorporating 3.6GB of packet buffers.

For more information on the HP 5920AF-24XG switch, please see hp.com/networking/5920.

The configuration for the HP 5920AF-24XG switch is provided below. The HP 5920 switch can also be used for connecting the SL4540 Shared Network Port LOM iLOs from the individual servers. (Direct connection of the iLO port on the SL4540 to the 5920 will not work, because the 5920 ports will not support 100Mbs.) This will require an HP X120 1G SFP RJ45 T Transceiver (JD089B) (1GbE transceiver) to be installed in the 5920 switch for each server. This connectivity option will also require that the iLO configuration on the servers be changed to utilize NIC 1, 1GbE port on the SL4540 server node (up to 3 per chassis), instead of the iLO module connector located in the SL4540 chassis. For additional information on how to configure the iLOs go to hp.com/go/ilo. Another option for iLO connectivity is to add an HP 5120 El switch and connect the iLO connectors from each chassis to that switch.

Note
For customers expecting a complete physical isolation between the iLO traffic from the servers as well as the management traffic, for their switches using the Out of Band interface (dedicated RJ-45 port), HP recommends to build a separate “out of band” network. By building this separate network, all the management traffic follows a completely different path which will never be impacted by any action executed on the data traffic coming from the servers (applications). A switch like the 5120 EI can be used for this type of “out of band network”, with 10/100/1000 RJ-45 connectivity (48 ports: JE069A or 24 ports: JE066A) to connect iLO ports.

Figure 6. HP 5920AF-24XG Top of Rack (ToR) switch

Capacity and sizing

One of the principal virtues of scale-out storage is that it can be scaled easily in multiple dimensions. In this solution, the flexibility of the HP ProLiant SL4540 Gen8 platform as well as that of Red Hat Storage Server makes multi-dimensional scalability a reality. In general, sizing constraints will be dictated by both the initial size of the data to be stored, as well as expectations for anticipated growth.

Workload description

To evaluate performance of the joint Red Hat and HP solution, testing was performed using a pair of two-node HP ProLiant SL4540 Gen8 servers along with four clients. Testing was conducted using IOzone—a filesystem benchmark tool that generates and measures a variety of file operations. IOzone is useful for performing a broad filesystem analysis of a particular platform and can characterize file system I/O based on a number of distinct operations.

Red Hat Storage Server configuration

For testing, two 2-node HP ProLiant SL4540 Gen8 servers were used to host Red Hat Storage Server. In this configuration, each HP ProLiant SL4540 Gen8 server contained two nodes, with each node controlling 25 disk drives for a total of four Red Hat Storage Server nodes in the storage cluster.

Each node in the cluster was equipped with two quad-core Intel Xeon E5-2407 processors, 48GB of Dual Rank x4 PC3L-10600R CAS-9 Low-Voltage memory, and a 10G IO Module Kit. For these tests, we utilized twenty-five 2TB 6G SAS 7.2K rpm LFF drives per storage node with two 500GB 6G SATA 7.2K drives in each node serving as the node operating environment. Each node within the HP ProLiant SL4540 Gen8 server was configured with an HP Smart Array B120i controller to manage the two 500GB drives serving the operating environment and an HP Smart Array P420i controller used to control the storage drives. Figure 7 illustrates the BIOS configuration of the B120i Smart Array controller.

Figure 7. BIOS for configuration of the HP Smart Array controller contained in the HP ProLiant SL4540 Gen8 server.

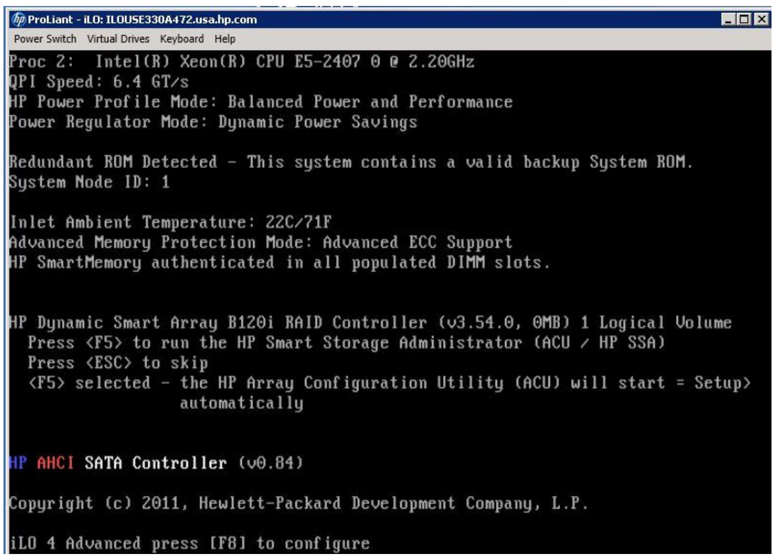

The HP Array Configuration Utility was then used to configure the 25 drives available on each node as follows (Figure 8):

- Two logical drives, each consisting of twelve 2TB SAS drives configured as an 18.2TB RAID 6 logical volume
- One unassigned drive

Figure 8. Each node within the 2-node HP ProLiant SL4540 server was configured with two 18.2TB RAID 6 logical volumes.

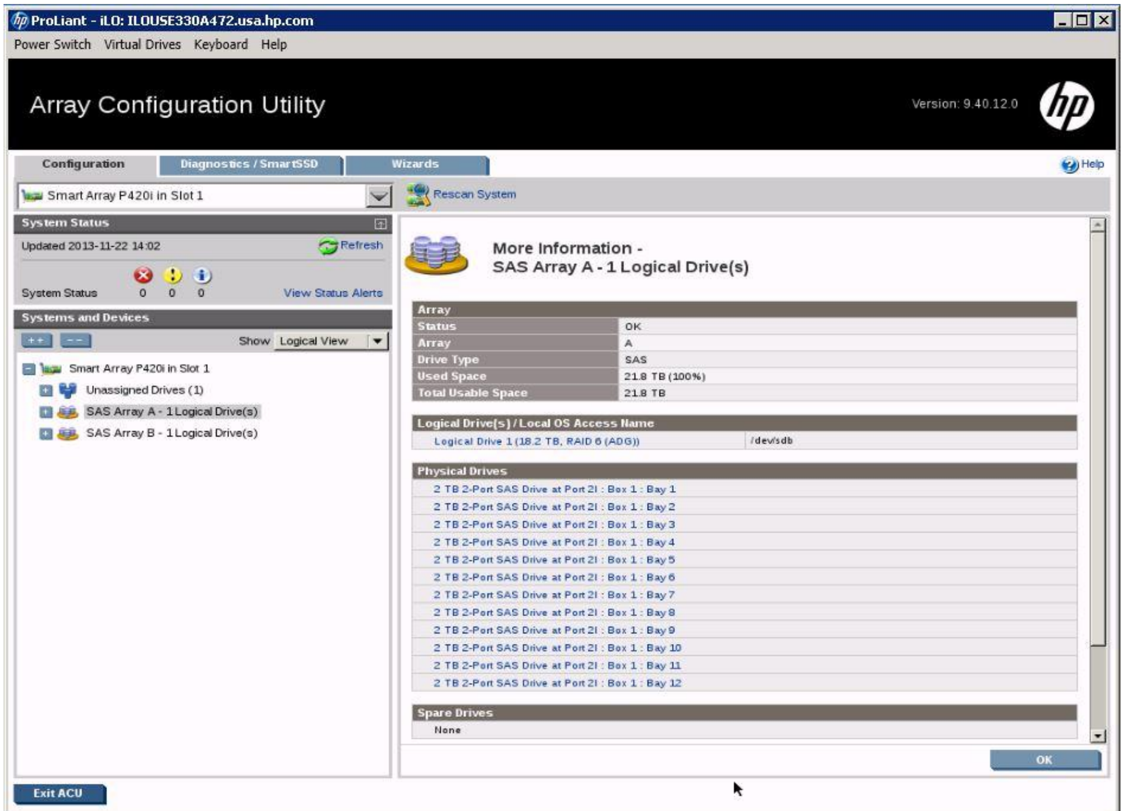

Red Hat Storage Server client configuration

For testing, four HP servers were used as clients. Each client system ran Red Hat Enterprise Linux 6.4 and had the Red Hat Storage Server client packages installed. Each client had two quad-core Intel Xeon E5-2407 processors, 96GB of Dual Rank x4 PC3L-10600R CAS-9 Low-Voltage memory, and a dual-port 10Gb Ethernet adapter.

Performance and scalability testing with IOzone

IOzone was used to test the performance of the reference architecture Red Hat Storage Server cluster. As mentioned, four Red Hat Storage Server nodes were configured, each was hosting a RAID6 based Red Hat Storage (RHS) brick and was running `glusterd`. Five Red Hat Storage Server volumes were created, including:

- Three Distribute volumes (distvol2, distvol3, distvol4)
- One Replicate volume (mirrorvol2)
- One Distributed-Replicate volume (mirrorvol4)

IOzone was run on up to four Red Hat Storage Server native clients, for I/O performance and scalability testing. IOzone was run in "Cluster testing" mode to launch eight IOzone tester threads on each client. With four clients, that capability resulted in the ability to test with 8, 16, 24, or 32 threads concurrently performing I/O to the Red Hat Storage Server cluster.

The IOzone command line is as follows:

```
rhs-client1# iozone -+m ${IOZONE_CONFIG_FILENAME} -i ${IOZONE_TEST}  -C  -w  -+n
-s ${IOZONE_FILESZ} -r ${IOZONE_RECORDSZ}  -+z  -c  -e  -t ${TEST_THREADS}
```

The following parameters were used in the IOzone command line:

- "-+m" specifies cluster testing mode
- IOZONE_CONFIG_ FILENAME is the IOzone config file for the cluster test format that lists the client hostnames and their respective GlusterFS mounts (that we created when mounting Red Hat Storage Server volumes).
- The IOZONE_TEST parameter was varied to cover the "Sequential Read", "Sequential Write", and "Random Read/Write" test cases.
- IOZONE_FILESZ was "8GB". 8GB file transfers were used as a representative "workload" for the "Content Cloud" reference architecture testing.

- IOZONE_RECORDSZ was varied between 64KB and 4MB, considering record sizes that were powers of two. This range of record sizes was meant to characterize the effect of record/block size (in client file requests) on I/O performance.
- TEST_THREADS specifies the number of threads or processes that are active during the measurement.

Workload results

Both distributed volumes and two-way replicated volumes were evaluated as a part of the testing. Both were tested for throughput as well as sequential and random reads and writes.

Test results—distributed volumes

Figure 9 illustrates the effect of record size (kiloBytes) variation, as defined by the IOzone record size variable, on RHS Distributed volume throughput (in MB/second). There are 4 clients with 8 threads each totaling 32 client threads targeting a 4-node Red Hat Storage Server cluster in a distribute volume configuration. The sequential read and sequential write saturate at record size greater than 64KB, and the random read and write scale over all reasonable record sizes.

Figure 9. The effect of request size on throughput for a distributed volume

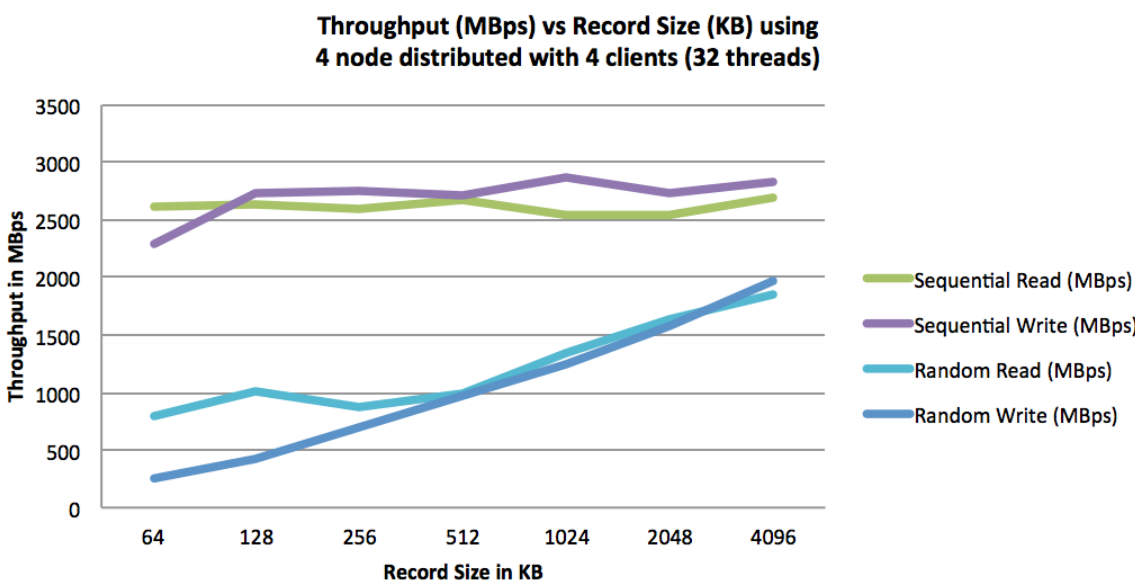

Figure 10 illustrates the effect of scaling out the number of Red Hat Storage Server nodes and clients on sequential writes for distributed volumes. The test used record size of 1024KB. With 8 client threads, the sequential write throughput is almost the same in spite of increasing the number of storage nodes. As we increase the total number of client threads from 8 (1 client) to 16 (2 clients) and then to 24 (3 clients), we see that the sequential write throughput increases significantly in each storage node configuration (2, 3 and 4 nodes). The increase in the sequential write throughput is less pronounced as we increase the total number of threads from 24 to 32 (4 clients).

Figure 10. Sequential write scalability for a distributed volume

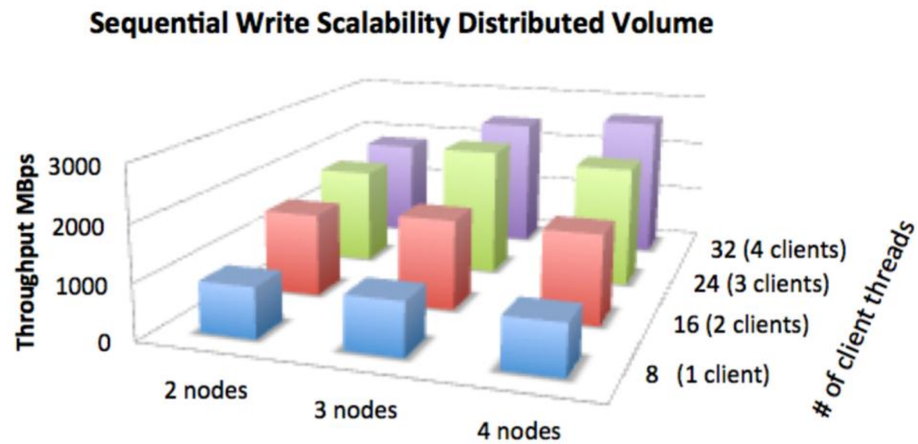

	2 nodes	3 nodes	4 nodes
8 (1 client)	905	950	897
16 (2 clients)	1511	1636	1623
24 (3 clients)	1773	2364	2216
32 (4 clients)	1842	2460	2651

Figure 11 illustrates the effect of scaling out the number of Red Hat Storage Server nodes and clients on sequential reads. The test used a record size of 1024KB. The sequential read throughput is almost the same in spite of increasing the number of storage nodes with 8 threads (1 client). As we increase the number of threads from 8 (1 client) to 32 threads (4 clients), we see that the read throughput increases appreciably with all the storage node configurations. The read throughput is higher with 3 nodes versus 2 nodes but almost the same for 3-node and 4-node configurations.

Figure 11. Sequential read scalability for a distributed volume

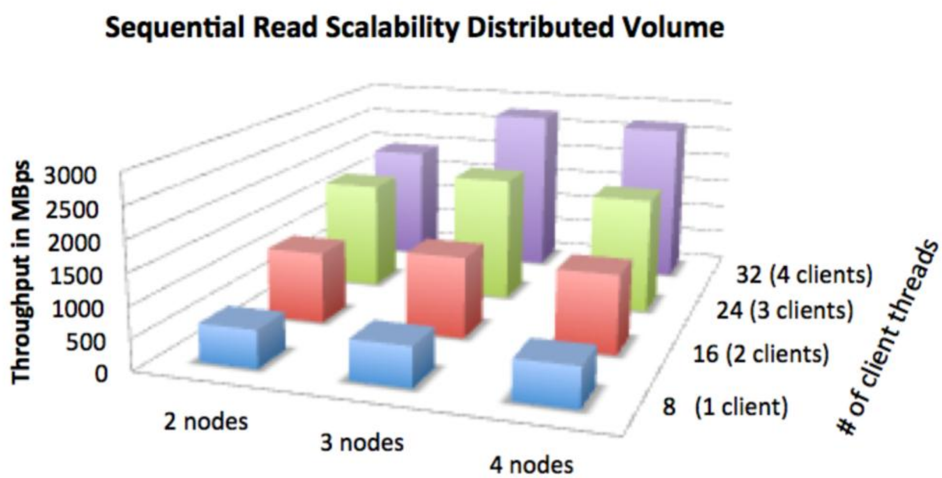

	2 nodes	3 nodes	4 nodes
8 (1 client)	605	635	631
16 (2 clients)	1164	1317	1274
24 (3 clients)	1785	2061	1902
32 (4 clients)	1953	2777	2683

Figure 12 illustrates the effect of scaling out the number of Red Hat Storage Server nodes and clients on random reads. The test used record size of 1024KB. The random read throughput does not increase appreciably with an increased number of threads in a 2-node configuration. Both in the 3-node and 4-node configurations, the random read throughput increases as the number of threads are increased from 8 (1 client) to 32 (4 clients). The read throughput in the 3-node and 4-node configurations are quite comparable.

Figure 12. Random read scalability for a distributed volume

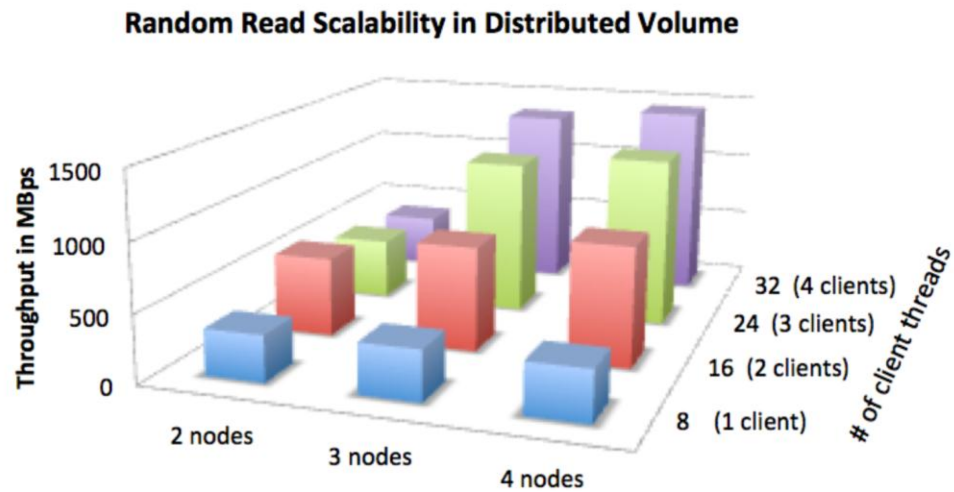

	2 nodes	3 nodes	4 nodes
8 (1 client)	340	369	370
16 (2 clients)	587	772	881
24 (3 clients)	451	1158	1261
32 (4 clients)	390	1348	1443

Figure 13 illustrates the effect of scaling out the number of Red Hat Storage Server nodes and clients on random writes. The test used record size of 1024KB. Increasing the number of client threads did not have an appreciable impact on the random write throughput in 2-node and 3-node configurations. In a 4-node configuration, the random write throughput increased significantly from 8 threads to 16 threads but remained flat with subsequent increments to 24 and 32 client threads.

Figure 13. Random write scalability for a distributed volume

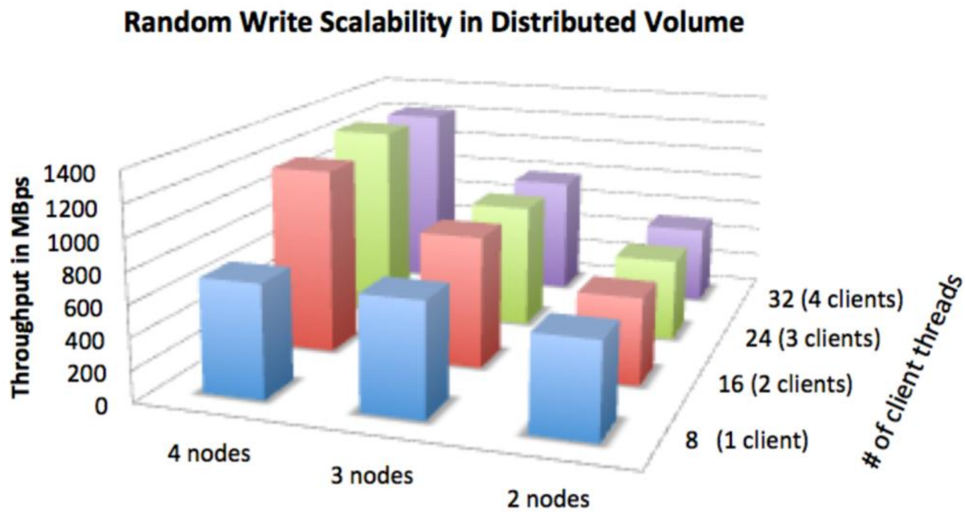

	4 nodes	3 nodes	2 nodes
8 (1 client)	718	717	596
16 (2 clients)	1188	842	554
24 (3 clients)	1265	812	532
32 (4 clients)	1226	785	516

Test results—two-way replicated volumes

Figure 14 illustrates the effect of request size variation (in kilobytes), as defined by the IOzone record size variable, on Red Hat Storage Server Distributed-Replicated volume throughput (MB/second). There are 4 clients with 8 threads each, totaling 32 client threads targeting a 4-node Red Hat Storage Server cluster in a distributed-replicate volume configuration. The sequential read and sequential write saturate at record size greater than 64KB. However, the random read and write performance scales over all reasonable record sizes.

Figure 14. Request size variation for distributed-replicated volumes

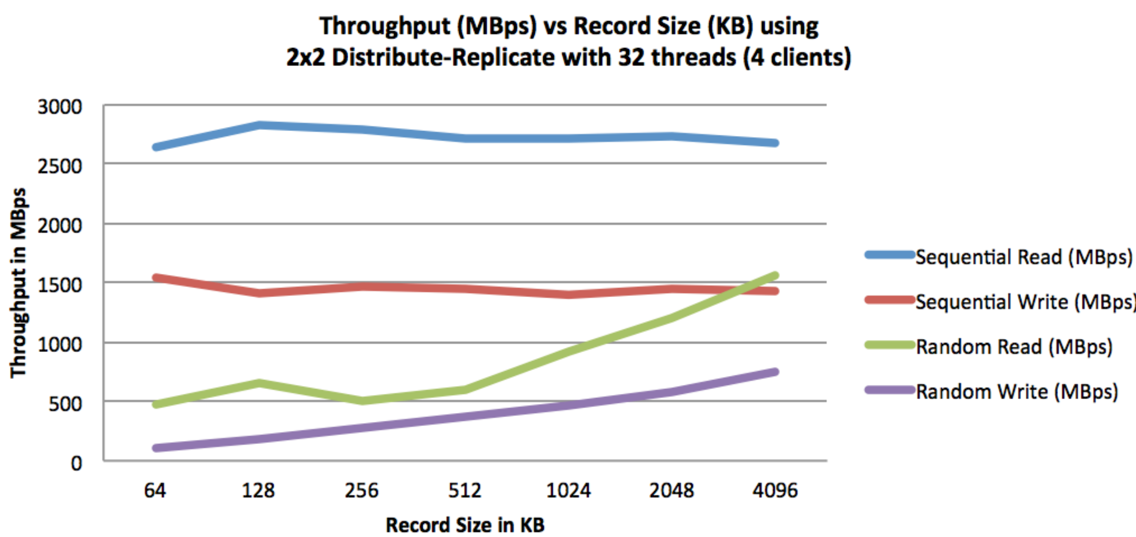

Figure 15 illustrates the effect of scaling out the number of Red Hat Storage Server nodes (in a distributed-replicated configuration) and clients on sequential writes. The test used a record size of 1024KB. The sequential write throughput shows a moderate increase from 8 to 16 client threads and then flat lines in the case of a 2-node replicated volume. The 4-node replicated volume doubles the sequential write throughput from 8 client threads to 16, and with subsequent increases in client threads the write throughput increases moderately. For 16 – 32 client threads workloads, the 4-node configuration shows twice the write throughput performance when compared to a 2-node configuration.

Figure 15. Sequential write scalability in 2-way replicated volume

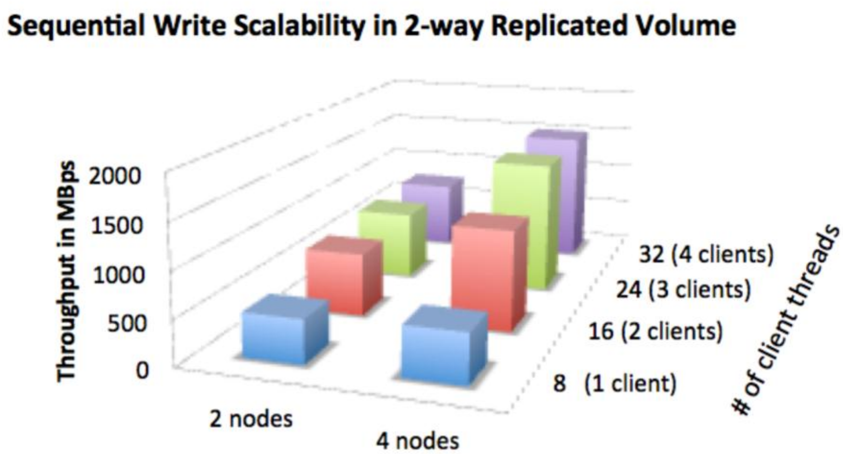

	2 nodes	4 nodes
8 (1 client)	485	548
16 (2 clients)	707	1129
24 (3 clients)	766	1499
32 (4 clients)	768	1518

Figure 16 illustrates the effect of scaling out the number of Red Hat Storage Server nodes (in a distributed-replicated configuration) and clients on sequential reads. The test used a record size of 1024KB. The sequential read throughput scales linearly as the Red Hat Storage Server nodes are added, albeit modestly. Sequential read throughput scales linearly as the number of client threads increases. Both in a 2-node and 4-node configuration, the sequential read throughput shows a linear increase from 8 to 32 client threads. The 4-node configuration shows a moderately higher sequential read throughput than a 2-node configuration.

Figure 16. Sequential read scalability in a 2-way replicated volume

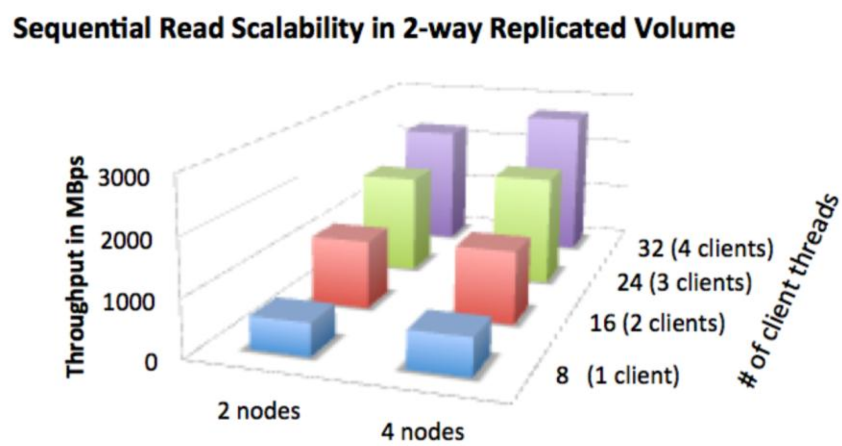

	2 nodes	4 nodes
8 (1 client)	566	647
16 (2 clients)	1216	1286
24 (3 clients)	1796	1979
32 (4 clients)	2252	2695

Figure 17 illustrates the effect of scaling out both the number of Red Hat Storage Server nodes (in a distributed-replicated configuration) and clients on random reads. The test used a record size of 1024KB. The 4-node configuration showed more than double the random read throughput than a 2-node configuration. Increasing the number of client threads had a detrimental effect on read throughput in a 2-node configuration. The random read throughput doubled when the number of client threads increased from 8 to 16 but peaked at 24 client threads.

Figure 17. Random read scalability in a 2-way replicated volume

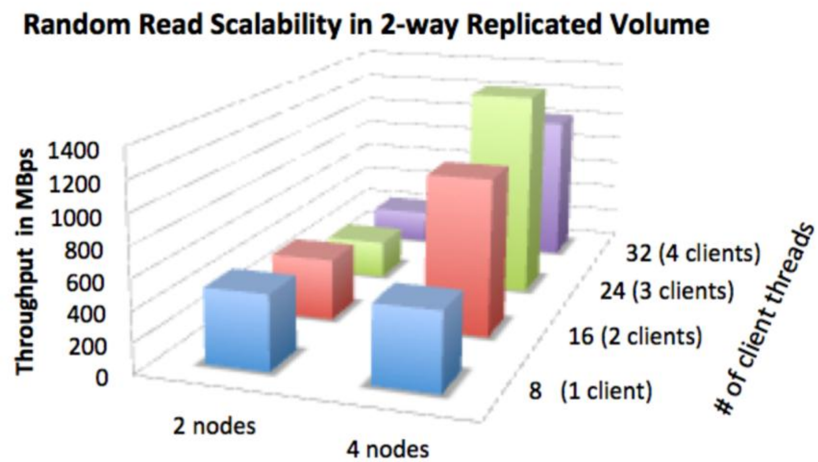

	2 nodes	4 nodes
8 (1 client)	492	511
16 (2 clients)	411	1046
24 (3 clients)	256	1394
32 (4 clients)	227	1021

Figure 18 illustrates the effect of scaling out both the number of Red Hat Storage Server nodes (in a distributed-replicated configuration) and clients on random writes. The test used a record size of 1024KB. The 4-node configuration shows double the random write throughput than a 2-node configuration. In case of both the 2-node and 4-node configuration, increasing the number of client threads does not increase the random write throughput.

Figure 18. Random write scalability in a 2-way replicated volume

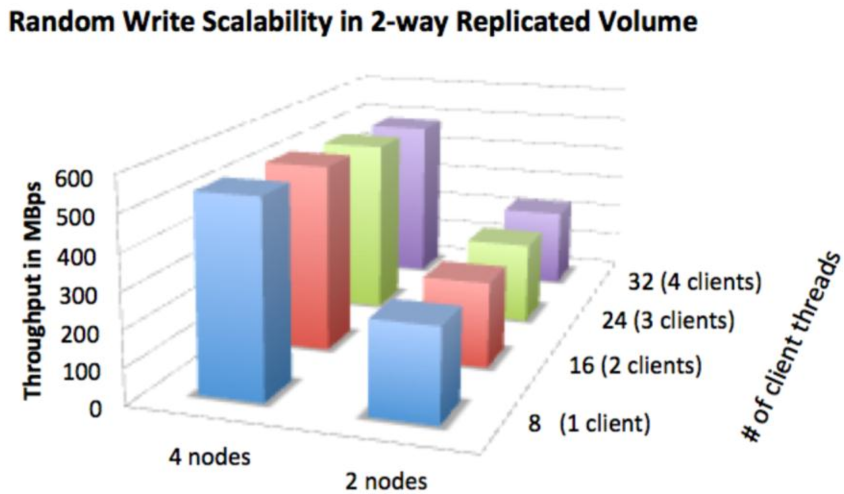

	4 nodes	2 nodes
8 (1 client)	539	255
16 (2 clients)	526	239
24 (3 clients)	501	232
32 (4 clients)	482	228

Analysis and recommendations

Scale-out storage using Red Hat Storage Server and HP ProLiant SL4540 Gen8 servers yields useful multi-dimensional scalability. Servers can be configured to ideally serve specific workloads and appropriate numbers of servers can then be deployed to match capacity needs. This building-block approach is particularly useful and allows organizations to scale storage infrastructure without arbitrary architectural limitations found in many competing solutions.

In general, scale-out storage deployments can be thought of as falling into small, medium, and large categories as described below. It is important to note that these are not fixed categories and any Red Hat Storage Server configuration can grow easily on demand.

- Small configurations typically comprise two-to-four HP ProLiant SL4540 Gen8 servers, for four-to-eight compute nodes and a total raw capacity of 200 to 800TB.
- Medium configurations span 16 to 20 nodes (8-10 two-node HP ProLiant SL4540 Gen8 servers). Depending on disks chosen, raw capacity can range from 800TB to 2PB.
- Large configurations can easily be built from 40+ nodes (20 or more two-node HP ProLiant SL4540 Gen8 servers), easily extending scale-out storage to multiple petabytes.

For larger configurations, a fully-configured rack of nine two-node HP ProLiant SL4540 Gen8 servers yields a particularly attractive hardware (and software) building block. With nine two-node servers, a full rack provides a raw capacity of from 450TB (when configured with 1TB drives) to 1.8PB (with 4TB disk drives). This range is a particularly useful and manageable size for a Red Hat Storage Server volume. Multiple racks and multiple Red Hat Storage Server volumes can be added as needed to construct scale-out storage of any size required.

One-, two-, and three-node HP ProLiant SL4540 Gen8 servers

HP understands that different applications and storage-related deployments require differing core-to-spindle ratios. The HP ProLiant SL4540 Gen8 server line offers different numbers of compute nodes, with different numbers of disk spindles supported by each compute node. While the use cases described in this document are ideal for the two-node HP server, Red Hat Storage Server can be run on any of the configurations as shown in Figure 19.

- One-node HP ProLiant SL4540 Gen8 Server. The one-node chassis supports a single compute node and as many as 60 attached disk drives. For the largest archival, object storage, or transient data applications with less need for compute performance scalability, the one-node HP ProLiant SL4540 Gen8 server enables maximum capacity scalability.
- Two-node HP ProLiant SL4540 Gen8 server. The two-node chassis supports more symmetrical application use cases, where some amount of computational power is required along with scalable data storage, density, and efficiency. Each of the two nodes in the chassis supports up to 25 drives. The large file and object store, enterprise drop box and cloud storage, and near-line archive use cases described herein are all ideal applications for the combination of the two-node HP ProLiant SL4540 Gen8 server and Red Hat Storage Server.
- Three-node HP ProLiant SL4540 Gen8 server. The three-node chassis provides an ideal platform for big data analytics applications such as Hadoop. With three compute nodes, each supporting up to 20 cores and up to 15 disk drives, the server exceeds the ideal 1:1 spindle to core ratio needed for the frequent interaction between processors and data dictated by many big data analytics applications.

Figure 19. One-node, two-node, and three-node HP ProLiant SL4540 Gen8 servers can all run Red Hat Storage Server to meet a range of application and capacity needs.

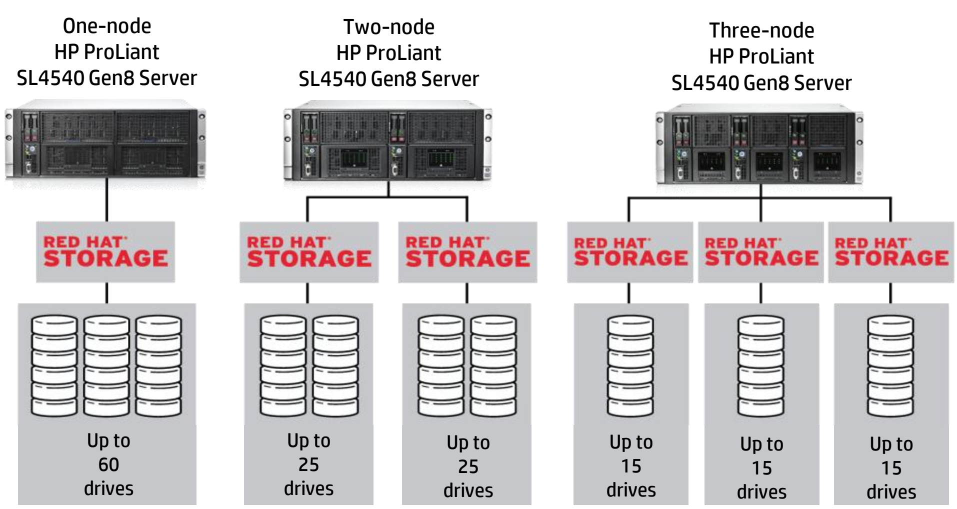

Server sizing for use cases

Individual HP ProLiant SL4540 Gen8 servers are widely configurable in terms of processors, memory, disks, and I/O to match specific workload requirements. While specific sizing recommendations are beyond the scope of this document, general recommendations can be made for configuring individual two-node HP ProLiant SL4540 Gen8 servers for use in scale-out storage configurations with Red Hat Storage Server. For all configurations, dual processors are recommended.

HP ProLiant SL4540 Gen8 servers support 4-, 6-, 8 -, and 10-core Intel E5-2400 and E5-2400 v2 series processors, depending on workload requirements. For the workloads described herein, processor performance is not typically a bottleneck. As a result, most deployments can deploy two quad-core Intel E5-2407 processors per node. Scenarios requiring higher computing requirements may wish to deploy processors with larger core counts.

- Large file and object store. For a large file and object store, HP and Red Hat recommend 48GB of DRAM for storage node system memory. Either 2TB SAS 7.2K rpm or 15K rpm drives are well suited for this application.
- Enterprise drop box or cloud storage. For enterprise drop box or cloud storage, HP and Red Hat recommend 48GB of DRAM for storage node system memory. Either 2TB or 3TB SAS 7.2K rpm or 15K rpm drives can be used for storage.
- Near-line archival. Near-line archival tends to be more focused on capacity than performance. For this reason, HP and Red Hat recommend 32GB of DRAM for storage node system memory. To maximize capacity, 3TB or 4TB SATA 7.2K rpm drives can be used for storage.

Configuration guidance

Red Hat Storage installation and configuration

All four nodes were installed with the RHS 2.1 ISO that is available at Red Hat Network (rhn.redhat.com), this release is based on RHEL 6.4 and GlusterFS 3.4. All the documentation for installing and configuring the Red Hat Storage Server can be found on the Red Hat website – https://access.redhat.com/site/documentation/en-US/Red_Hat_Storage/ (see Figure 20).

Figure 20. Red Hat Storage website

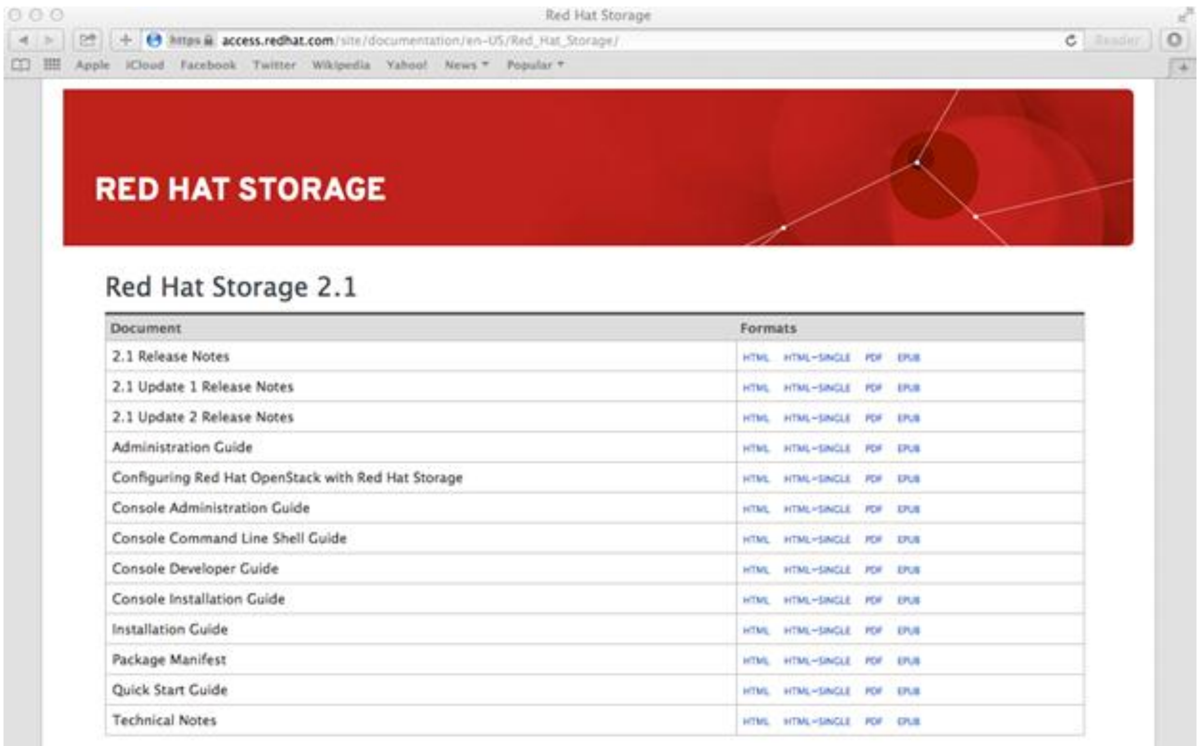

Install the Red Hat Storage Server by following the instructions in the Installation Guide. Once the system is up and running, you will create the storage bricks using the rhs-server-init.sh script. This script does the following:

- Create physical volume
- Create Logical volume
- Make XFS file system on the logical volume
- Run tuned performance profile for high-throughput

Red Hat Storage client software installation

After you have successfully installed your client operating system, you must first register the target system to the Red Hat Network and subscribe to the Red Hat Enterprise Linux Server channel.

To subscribe to the Red Hat Enterprise Linux Server channel using RHN Classic:

1. Register the system
   Run the # `rhn_register` command to register the system with Red Hat Network. To complete registration successfully you will need to supply your Red Hat Network username and password.

   In the `select operating system release` page, select `All available updates` and follow the prompts to complete the registration of the system. The system is now registered to `rhel-x86_64-server-6` channel.

2. Subscribe to Red Hat Storage Native Client.
   You must subscribe the system to the Red Hat Storage Native Client channel using either the web interface to Red Hat Network or the command line rhn-channel command.

   A. Using the **`rhn-channel`** command.
      You can run the rhn-channel command to subscribe the system to the Red Hat Storage Native Client channel. If your client version is Red Hat Enterprise Linux 6.x, then run the following command to subscribe:
      ```
      # rhn-channel --add --channel=rhel-x86_64-server-rhsclient-6
      ```
      If your client version is Red Hat Enterprise Linux 5.x, then run the following command to subscribe:
      ```
      # rhn-channel --add --channel=rhel-x86_64-server-rhsclient-5
      ```

   B. Using the Web Interface to Red Hat Network.
      To add a channel subscription to a system from the web interface:

      i. Log on to Red Hat Network (rhn.redhat.com).
      ii. Move the mouse cursor over the Subscriptions link at the top of the screen, and then click the Registered Systems link in the menu that appears.
      iii. Select the system to which you are adding Red Hat Storage Native Client channel, from the list presented on the screen, by clicking the name of the system.
      iv. Click Alter Channel Subscriptions in the Subscribed Channels section of the screen.
      v. On this screen, expand the node for Additional Services Channels for Red Hat Enterprise Linux 6 for x86_64 for RHEL 6 or Additional Services Channels for Red Hat Enterprise Linux 5 for x86_64 for RHEL 5.
      vi. Click the Change Subscriptions button to finalize the changes. After the page refreshes, select the Details tab to verify if your system is subscribed to the appropriate channels.

      Run the following command to verify if the system is registered successfully.
      ```
      # rhn-channel –l
      rhel-x86_64-server-6
      rhel-x86_64-server-rhsclient-6
      ```

3. The system is now registered with Red Hat Network and subscribed to the Red Hat Storage Native Client channel. Now install the native client RPMs using the following command:
   ```
   # yum install glusterfs glusterfs-fuse
   ```

Red Hat Storage volume creation

Once you have created and mounted bricks on all RHS server nodes, you will need to instantiate a few RHS volumes that are useful for recreating the test setup we describe here. We will create “distribute” and “replicate” (mirrored) flavors of RHS volumes, specifically:

1. 2, 3, 4-node configurations of Distribute Volumes
2. A 2-node Replicate Volume
3. A 4-node Distribute-Replicate Volume

Before we do, make sure that you have “`glusterd`” daemon running on each of the Red Hat Storage servers (rhs01, rhs02, rhs03 and rhs04). Using the gluster peer probe command create the storage cluster from rhs01 server:

```
rhs01 # gluster peer probe rhs02

rhs01 # gluster peer probe rhs03

rhs01 # gluster peer probe rhs04
```

Confirm that all the storage servers are in connected state using the peer status command:

```
rhs01 # gluster peer status
```

A. Create Distribute only volume:

```
rhs01 #  gluster volume create  dist2  rhs01:/rhs/brick1/dist2
rhs02:/rhs/brick1/dist2

rhs01 #  gluster volume create  dist3  rhs01:/rhs/brick1/dist3
rhs02:/rhs/brick1/dist3  rhs03:/rhs/brick1/dist3

rhs01 #  gluster volume create  dist4  rhs01:/rhs/brick1/dist4
rhs02:/rhs/brick1/dist4  rhs03:/rhs/brick1/dist4
rhs04:/rhs/brick1/dist4
```

B. Create Replicate volume:

```
rhs01 #  gluster volume create  mirror2 replica 2
rhs01:/rhs/brick1/mirror2  rhs02:/rhs/brick1/mirror2
```

C. Create Distribute-Replicate volume:

```
rhs01 #  gluster volume create  mirror4 replica 2
rhs01:/rhs/brick1/mirror4  rhs02:/rhs/brick1/mirror4
rhs03:/rhs/brick1/mirror4    rhs04:/rhs/brick1/mirror4
```

Mount RHS volumes on client nodes

Mount all the RHS volumes that you created on all the four clients using the following commands:

```
client01 # mkdir /mnt/dist2  /mnt/dist3  /mnt/dist4  /mnt/mirror2  /mnt/mirror4
client01 # mount  -t  glusterfs  rhs01:/dist2  /mnt/dist2
client01 # mount  -t  glusterfs  rhs01:/dist3  /mnt/dist3
client01 # mount  -t  glusterfs  rhs01:/dist4  /mnt/dist4
client01 # mount  -t  glusterfs  rhs01:/mirror2  /mnt/mirror2
client01 # mount  -t  glusterfs  rhs01:/mirror4  /mnt/mirror4
```

Repeat these commands on the other three clients.

Note
Even though we are using the same RHS server for mounting the volumes, the Native protocol has built-in load balancing. The clients use the mount server initially to get the volume information, and after that they will contact the individual storage servers directly for accessing the data. The data request does not have to go through the mount server.

Bill of materials

The bill of materials found in Table 1 lists the major server and storage hardware components needed for the reference architecture, however, this is not an exhaustive listing of all the necessary components needed to build the complete solution. For complete configuration details, please contact your HP Reseller or HP Sales Representative.

This bill of materials is for a sample small configuration of four two-node storage servers (eight nodes). Due to the modular, scale-out approach to this solution, organizations can scale this solution up or down depending on their required capacity.

Note
Part numbers are at time of publication and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HP Reseller or HP Sales Representative for more details. hp.com/large/contact/enterprise/index.html

Table 1. Bill of materials

Qty	Description	Part Number
4	HP SL4500 2 Node Chassis	663600-B22
8	HP SL4540 Gen8 Tray 2 Node Server	664644-B22
8	HP SL4500 10G IO Module Kit	664648-B21
8	HP Smart Array P420i Mezz Ctrllr FIO Kit	692276-B21
8	HP SL4500 Storage Mezz to PCIe Opt Kit	682632-B21
8	HP 12in Super Cap for Smart Array	668943-B21
8	HP 2GB FBWC for P-Series Smart Array	631681-B21
16	HP 750W CS Plat PL Ht Plg Pwr Supply Kit	656363-B21
4	HP 4.3U Rail Kit	681254-B21
1	HP 0.66U Spacer Blank Kit	681260-B21
1	HP 13U FIO Rack for 3X4.3U Adapter Kit	681677-B21
8	HP SL4540 Gen8 E5-2407 FIO Kit	686166-L21
8	HP SL4540 Gen8 E5-2407 Kit	686166-B21
48	HP 8GB 2Rx4 PC3L-10600R-9 Kit	647897-B21
16	HP 500GB 6G SATA 7.2k 2.5in SC MDL HDD	655708-B21
200	HP 2TB 6G SAS 7.2K 3.5in SC MDL HDD	652757-B21
1	HP 59xx CTO Switch Solution	JG505A
1	HP 5920AF-24XG Switch	JG296A
24	HP X130 10G SFP+ LC SR Transceiver	JD092B
2	HP A58x0AF 650W AC Power Supply	JC680A
2	JmpCbl-NA/JP/TW	JC680A-B2B
2	HP 5920AF-24XG Bkpwr-Frtprt Fn Tray	JG297A

Summary

With the combination of Red Hat Storage Server and the HP ProLiant SL4540 Gen8 server, scale-out storage is now a viable solution for large amounts of structured and semi-structured data. As a software-defined and open source solution, Red Hat Storage Server brings significant advantages to a range of use cases, from large file and object stores, to enterprise drop boxes and cloud storage, to near-line archival. With a choice of processors, memory, disks, and I/O, HP ProLiant SL4540 Gen8 servers represent a modular and scalable high-capacity building block for deploying scale-out storage. Together, HP and Red Hat let organizations deploy software-defined scale-out storage that fits both the characteristics of their applications and their escalating data storage needs.

Implementing a proof-of-concept

As a matter of best practice for all deployments, HP recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HP Services representative (hp.com/large/contact/enterprise/index.html) or your HP partner.

For more information

To read more about the HP/Red Hat Alliance and HP solutions for Red Hat, please see: hp.com/go/redhat

For additional information on the HP ProLiant SL4540 Gen8 Server, please see: hp.com/go/proliant/bigdataserver

Additional information about Red Hat Storage can be found at: redhat.com/storage. You can watch a Red Hat Storage Demo at redhat.com/promo/liberate/2222_RedHat_Storage.html, and be sure to engage Red Hat to begin a Proof Of Concept: https://engage.redhat.com/forms/storage-poc

Additional information on the IOzone Filesystem Benchmark test used in this reference architecture can be found at: iozone.org

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.

Sign up for updates
hp.com/go/getupdated

4AA4-7192ENW, May 2014